March-April 1997
Vol. 2, No. 23

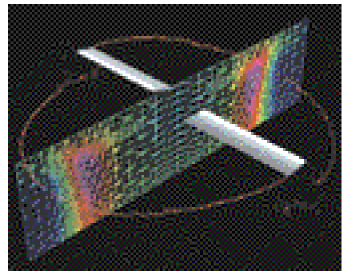

- Vortex Core Detection Technique Helps Rotorcraft Research
- Onward to Petaflops: Supercomputing for the 21st Century
- New 'NAS' Mission: Stick to High-risk Pathfinding
- New Tool For Analyzing Wind Tunnel Experiment Data
- New Highbatch Queue on the CRAY C90s
- Filesystem Reconfiguration Makes Sense for ACSF Users
- 'Commodity' Cluster Prototype Ready for Kickoff
- Lewis Aeroacoustic Study Reveals Jet Noise Patterns
- Changes Made To CRAY J90 System
- UIG Meeting Brings Users Up To Date
- Credits

Vortex Core Detection Technique Helps Rotorcraft Research

by David Kenwright

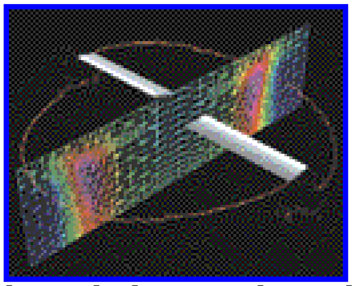

Ames Research Center is the NASA-designated Center of Excellence for both rotorcraft research and information technologies. Scientists from both disciplines have teamed up to solve one of the most challenging problems in rotorcraft computational fluid dynamics (CFD) accurate prediction of the rotor wake.

The rotor wake is the disturbed flow that is left behind as a rotor blade cuts through the air. The wake rolls up into vortices near the blade tips due to the pressure differences caused by the moving rotor. Vortices are usually an unwanted side-effect of lift, since they cause control buffeting, vibration, and audible noise. The requirement for low noise is particularly important for civilian helicopters and tiltrotors, which operate in highly populated areas.

Problem With Diffusion

A major problem with CFD simulations of rotorcraft is that the tip vortices diffuse too quickly due to inadequate grid resolution. Adaptive grid refinement schemes developed at Ames by Rupak Biswas, Data Analysis Branch, and Roger Strawn, Applied Aerodynamics Computational Branch (AAC), have been partially successful in capturing these vortice -- although inadequate grid resolution around the vortex core still causes discrepancies from experimental results.

A promising technique to help solve the vortex diffusion problem comes from a scientific visualization technique that locates the vortex cores in CFD datasets. This technique was recently implemented in UFAT (the Unsteady Flow Analysis Toolkit) as part of a joint research project between members of the U.S. Army Aeroflightdynamics Directorate and the data analysis group in the NAS Systems Division.

The diffusion problem will be tackled by linking the vortex core extraction code with an unstructured grid adaption code, currently being developed by Biswas and Strawn. The vortex cores will be used to guide the grid adaption code by identifying the regions of the grid that need refining or coarsening. This process will occur automatically during the flow computations.

Vortex Cores Steer Grid Generation

Grid generation, flow simulation, and flow visualization are usually treated as separate tasks. In this collaborative project, the loop between them is closed by using the vortex core extraction technique to guide the grid generation process. It is believed that this approach will reduce computation times and lead to more accurate simulations of rotorcraft.

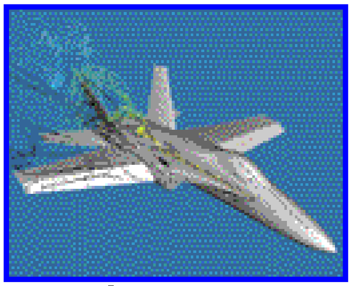

The figure at left shows a typical vortex core trace for a hovering rotor. The solution was computed by Biswas and Strawn with an unstructured grid Euler solver developed by Tim Barth, Ames Office of Aeronautics.

The vortex core detection algorithm implemented in UFAT was based on an eigenvector method developed by Robert Haimes and David Sujudi of the Massachusetts Institute of Technology. The initial research was jointly funded by Ames and United Technologies Research Center. Continued development of the eigenvector method at Ames has led to improved performance and accuracy of the algorithm. The implementation in UFAT also permits the time evolution of vortex cores to be studied in unsteady flows.

Vortex cores are extracted in a batch computation in UFAT, which can take anywhere from seconds to hours depending on the size of the dataset and the complexity of the flow. For the unstructured grid (shown at top), which contains approximately 800,000 tetrahedral elements, the extraction process took only eight seconds on a Silicon Graphics Onyx workstation with one R10000 processor.

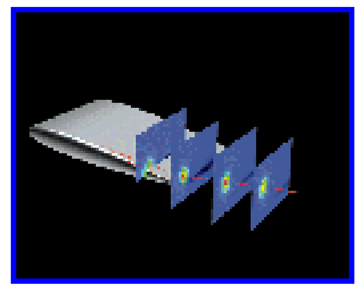

In contrast, it took six hours of computation on the same computer to extract the vortex cores from the F-18 aircraft dataset (shown in the figure at right). That dataset had over 1.5 million grid points and 285 solution fields. Ken Gee (AAC), who ran the flow simulation for the F-18, remarked how useful this technique was to identify the location of the vortex breakdown.

Comparing Numerical & Experimental

Wind tunnel data has also been analyzed by the vortex core algorithm. The bottom figure shows results from a numerical/experimental study of a wing tip vortex by Jennifer Dacles-Mariani and Greg Zilliac, both at Ames. The vortex cores extracted from both experimental and simulated datasets were overlaid to validate the measured error of 3 percent. UFAT users can invoke the vortex core extraction algorithm using the *-vc* command-line option. The results from the program are written into an ARCGraph file that can be displayed using FAST (Flow Analysis Software Toolkit).

For additional information on the vortex core extraction technique, send email to davidk@nas.nasa.gov.

David Kenwright is a research scientist working in the NAS Systems Division's Data Analysis Branch. He is currently developing feature identification techniques that provide automated analysis of large numerical datasets.

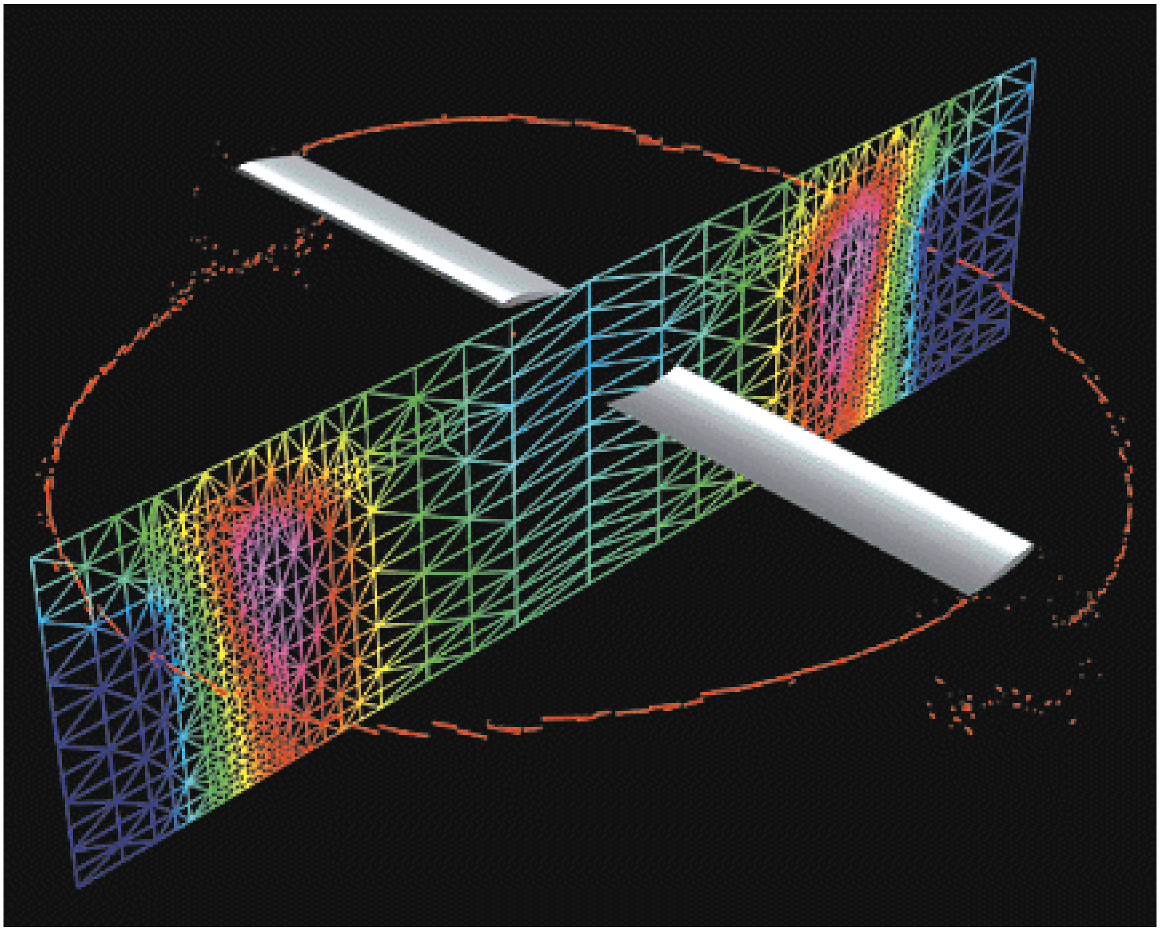

Vortex cores identify the center of swirling or rotating flow. The vortex cores shown here (in red) were extracted from a simulation of a helicopter rotor in hover. The vortex cores, which originate at the rotor tips, progressively diffuse as they are left behind in the rotor wake-- indicated by the disjointed line segments. Researchers at NASA Ames Research Center are collaborating on visualization techniques to accurately predict rotor wake. Image by David Kenwright, NAS data analysis group; flow computation by Rupak Biswas and Roger Strawn, both at NASA Ames.

to the article.

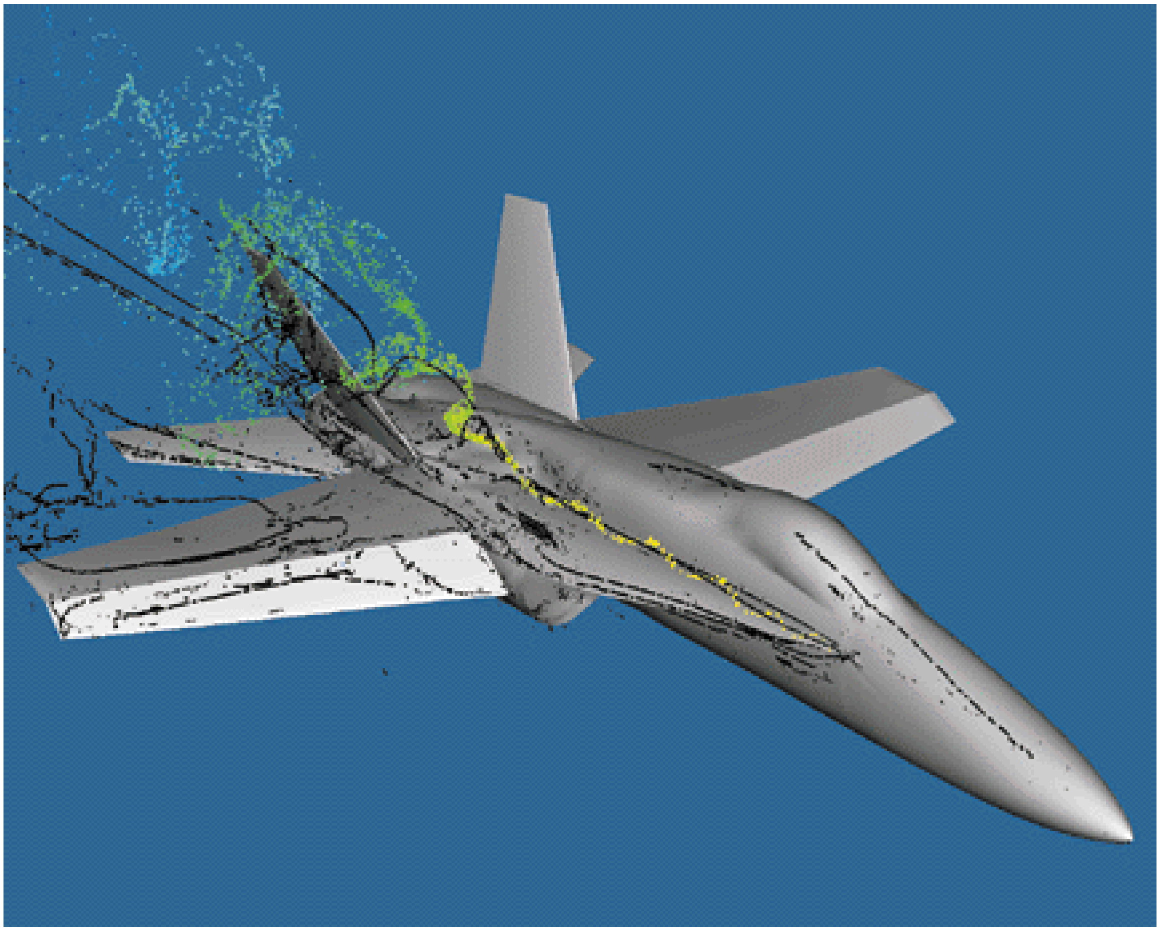

Above, a frame from an animation of vortex cores around an F-18 aircraft. The vortex cores (black lines) were extracted from 204 timesteps of data in a batch computation using UFAT, the Unsteady Flow Analysis Toolkit. Colored streaklines wind around the vortex of primary interest, which was known to cause tail buffeting.

Below, the vortex cores (shown in red) were used to compare results from wind tunnel experiments with numerical (computational fluid dynamics) data. The vortex cores identify the precise center of swirling flow -- often difficult to establish using the blue cutting planes.

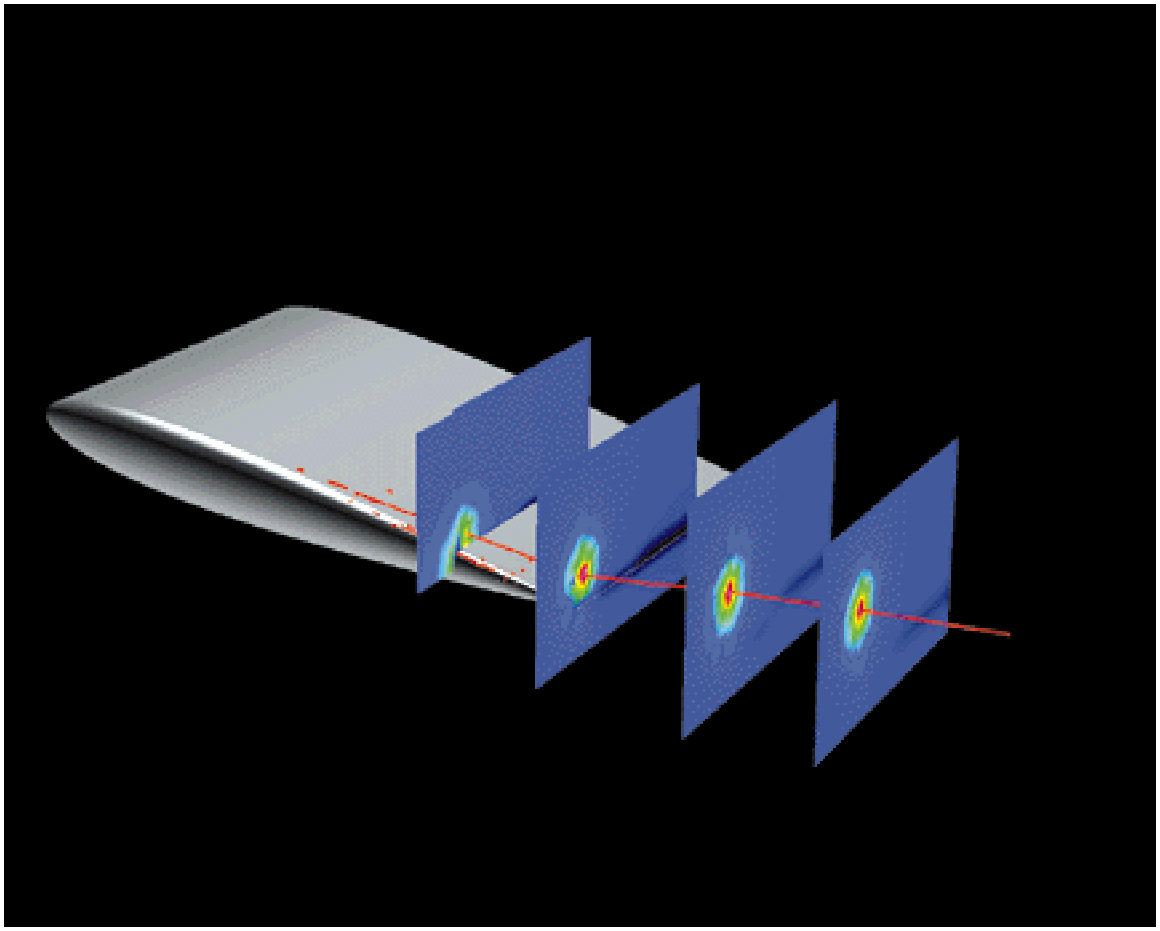

Images by David Kenwright.

 to the article.

Onward to Petaflops:

Supercomputing for the 21st Century

by David H. Bailey

With programs such as the U.S. High Performance Computing and Communications (HPCC) Program, the attention of scientists and engineers worldwide has been focused on the potential of very high performance scientific computing -- systems that are hundreds or thousands of times more powerful than those typically available in desktop systems at any given point in time. Extending the frontiers of computing in this manner has resulted in remarkable advances, both in computing technology and in the various scientific and engineering disciplines that utilize these systems.

On December 16, a sustained rate of 1 Tflop/s (1 teraflops, or 10^12 floating-point operations per second) was achieved by the "ASCI Red" system, which employs some 7,000 Intel Pentium Pro processors at Sandia National Laboratory in New Mexico. Now that this long-sought milestone has been achieved, what lies ahead for high-end computing?

The next major milestone is a sustained rate of 1 Pflop/s (1 petaflops, or 10^15 floating-point operations per second). We could just as well use the term "peta-ops," since it appears that large scientific systems will be required to perform intensive integer and logical computation in addition to floating-point operations, and completely non-floating-point applications are likely to be important as well.

In addition to prodigiously high computational performance, such systems must of necessity feature very large main memories -- between 10 Tbyte (10^13 byte) and 1 Pbyte (10^15 byte), depending on application -- as well as commensurate I/O bandwidth and huge mass storage facilities. The current consensus of scientists who have performed initial studies in this field is that "affordable" petaflops systems may be feasible by the year 2010, assuming that certain key technologies continue to progress at current rates.

To get some idea of the scale of these systems, a 1-Pflop/s computer could dispatch in three seconds a computation that a current desktop workstation would require a full year to perform. One Pbyte of memory could contain the text of approximately one billion books -- roughly 1,000 times the size of a typical university library. If a petaflops system were constructed today using low-cost personal computer components (ignoring for a moment the daunting difficulties of communication and software for such a system), it would cost some $50 billion and would consume about 1,000 megawatts of electric power.

History Shows Need for Petaflops

The need for such enormous computing capability is often questioned, but such doubts can be dismissed by a moment's reflection on the history of computing. It is well known that Thomas J. Watson, a founder of IBM, once ventured that there was a worldwide market of only about six computers. Even the legendary Seymour Cray (who recently passed away) designed his CRAY-1 system on the premise that there were only about 100 potential customers. In 1980, after the CRAY-1 had already achieved significant success, an internal IBM study concluded that there was a limited market for supercomputers, and as a result IBM delayed its entry into the market.

In stark contrast to these short-sighted projections, some private homes now have more than Watson's predicted six systems. Further, the latest personal computers have more computational power and main memory than the original CRAY-1, and enthusiastic users are clamoring for more. When we observe home-based users becoming annoyed with the "slowness" of their Cray-class computer when running a spreadsheet or graphics rendering application, it is easier to understand how enormous physical simulations could strain the computing power of existing supercomputers.

Demand Likely to Increase

High-end scientific computers traditionally have been the province of academic and government research laboratories. But in a significant recent development, parallel supercomputers are increasingly being used by professionals in other arenas, including financial analysts in the Wall Street community and marketing analysts in the consumer banking and retail industries. These developments are certain to increase the demand for future high-end computers.

Since the demand for state-of-the-art computing power appears insatiable, we may as well start planning now for petaflops systems. Some of the compelling applications anticipated for petaflops computers include the following:

- Nuclear weapons stewardship
- Satellite data processing
- Climate and environmental modeling
- 3-D protein molecule reconstructions
- Severe storm forecasting
- Design of advanced aircraft
- Molecular nanotechnology
- Intelligent planetary spacecraft

To elaborate on just a single item, consider 3-D protein molecule reconstructions, also known as the "protein folding problem." In designing a new drug agent, scientists need to examine many protein molecules, each with a specified nucleotide sequence. But at present it is not possible to reliably determine, except by laboratory experiment, the actual 3-D structure of the resulting protein molecule. And without this knowledge, it is not possible to know whether the molecule will have the proper binding sites to be an effective drug agent. Petaflops computers may be powerful enough to do the necessary computations to determine this 3-D structure in a reasonable amount of time. Such a capability could be a powerful new tool for pharmaceutical research.

Some of these anticipated petaflops computer applications will be scaled-up versions of present-day applications, with evolutionary enhancements. Others will consist of integrated simulations of multiple physical effects. Many of these applications will likely employ advanced visualization facilities, such as immersive or remote visualization environments, which are still under development. But if the history of computing is any guide, a number of exotic new applications will be enabled by petaflops computing technology. These applications may have no clear antecedent in today's scientific computing, and in fact may be only dimly envisioned at present.

Government Funding Will Be Needed

In spite of such potential, it is by no means certain that scientific computers produced by private industry will achieve the level of 1 Pflop/s by 2010. One reason for this conclusion is the recent turmoil in the scientific computing marketplace, which has led computer vendors to cut long-term research in favor of near-term development, and to focus on the more lucrative low- and mid-level systems instead of high-end systems. This phenomenon has been described as the "truncated pyramid" of the current computing marketplace. In this environment, it is likely that government agencies will need to provide a substantial part of the funding for the research and development needed to make these systems a reality.

A number of difficult technical problems need to be solved in the next few years in order to achieve the goal of petaflops computers by the year 2010. Indeed, the anticipated difficulties of developing the necessary hardware technology, determining an optimal system architecture, producing reliable system software, devising efficient algorithms, and ultimately programming petaflops systems, present challenges unprecedented in the history of computing.

Latency: No `Dramatic Improvements'

A key issue for these systems is latency management. When citing the breathtaking increases in memory device density during recent years, a consequence of Moore's law, it is often forgotten that the access time of these memory devices has not improved much during this time -- nor is there any reason to expect dramatic improvements in the foreseeable future. Thus the gap between processor speed and memory speed is expected to worsen in the future.

Latency can been dealt with by exploiting concurrency, such as in pipelined or multi-threaded architectures. Overcoming latency, coupled with the need to achieve 1 Pflop/s aggregate sustained performance, means that enormous system concurrency will be required -- possibly up to 1,000,000 processors. Concurrency of this scale is well beyond anything heretofore attempted in high performance computing. Indeed, the dual challenges of managing latency and extreme concurrency will drive much of the research that needs to be done over the next few years.

A growing research community is already working on these and related problems of petaflops computing. For example, the National Science Foundation recently awarded a number of research grants to explore system architectures for petaflops computers. These researchers presented their preliminary results at the Frontiers '96 conference, held in Annapolis, MD last October. More studies are planned. Onward to petaflops computing!

Intriguing Questions For Petaflops Research

Some specific research questions that need to be answered include the following:

Hardware

- Can we produce a usable petaflops system using commercial, off-the-shelf (COTS) hardware components?
- Is a hybrid hardware technology approach, such as superconducting RSFQ logic with optical interconnect, superior to a COTS-based design?
- Can the power consumption for such a system be reduced to acceptable levels?
- Will a multiple-instruction, multiple-data (MIMD) distributed memory architecture be satisfactory, or will some novel system architecture be required?
- What hardware facilities are needed to manage latency and multiple layers of memory hierarchy?

Software

- What operating system design can reliably manage 100,000 to 1,000,000 processors?
- Are radically new programming languages needed, or can existing languages be extended?
- What specific new language constructs will be required?
- What is the best way to support I/O, debugging, and visualization?

Algorithms

- Do latency-tolerant variants of known algorithms exist?
- How will the operation count, memory requirement, data locality and other characteristics of various algorithms scale on these future systems?
- Will variations of classic algorithms suffice for key applications, or will completely new algorithms be needed?

Applications

- Can anticipated petaflops applications be structured to exhibit the required 100,000-plus concurrent threads?

- What is the best way to implement various applications on proposed system designs?

- What will be the memory and I/O requirements of future applications? What exotic new applications might be enabled by petaflops systems?

These research questions raise provocative issues about the future of all computing, not just high-end scientific computing. After all, high levels of parallelism are inevitable for all classes of computing, even desktop systems. Thus, it is likely that answers to these questions may have impact far beyond the realm of large-scale scientific computing.

to the article.

New 'NAS' Mission: Stick to High-risk Pathfinding, Hand Off Production Computing Management

by Elisabeth Wechsler

The NAS Program, one of three funding components of the NAS Systems Division, has been re-chartered as a pathfinder in integrated supercomputing systems that address "tremendous requirements in computing, data storage, networks, systems software, and intelligent aids for decision making," according to Tom Edwards, program manager for Information Technology (IT). The former NAS Program -- officially named Advanced Computing, Networks, and Storage (ACNS) -- is now part of IT. *(See detailed explanation.)*

"Once again we're going to be out on the 'bleeding edge' of pathfinding," Edwards said. At the same time, ACNS will hand off its role as a production computing provider to CoSMO, NASA's Consolidated Supercomputing Management Office. "If you're operating a production system, you're not on the leading edge, so we're going to be investing in systems that in some cases may not prove out their potential but, on the other hand, may provide significant breakthroughs."

As he sees it, the former NAS Program may have been a victim of its own success: "For the first 10 years or so, CFD was really the driving application that made the program what it is today. We succeeded but then stopped pushing the frontier."

Edwards believes that part of the reason the NAS Program strayed from advancing state-of-the-art supercomputing in the past several years was because vonneumann (the CRAY C90), at almost four years old, is considered mature technology. "It's become a very stable platform and hence not as useful for supercomputing research," he said. "We're going to reinforce the original NAS vision, which is to invest in higher-risk research that has a higher payoff."

Integrated and Interactive Capabilities

Edwards doesn't shy away from questions about future challenges. In general, he views ACNS's top priorities as those which emphasize system integration and interactive capabilities. For starters, he wants increased focus on supercomputing systems to support integrated design systems that combine CFD with the engineering disciplines needed for solving design problems. He wants systems that will integrate real-time data -- for example, from wind tunnel tests -- with computational results to visualize it and to conduct "very productive, efficient tests on the ground." He pointed to the DARWIN project as an "excellent example of an integrated design system project."

Another emerging application being evaluated by IT is an aviation "internet" of sorts which would bring together all information systems that support aviation and integrate them into a "system of systems," allowing rapid decision-making about routing airplanes and other safety concerns. The overall goal is to simultaneously improve aviation safety while increasing the airspace capacity so that more planes can fly.

As in the past, space applications and other non-aeronautics applications will continue to play a role in advancing system performance. They "tend to push the state of the art fairly hard -- and we can get a lot of benefit from them in learning what the system capabilities and requirements need to be," Edwards said. "We're going to explore some fundamental technologies that will enable revolutionary increases in supercomputer performance. We're aiming at petaflops computing with this technology, and to get there we need to work fundamental issues such as device modeling, nanotechnology, and low-power, lightweight systems."

For example,"A lot of space probes now are relatively unintelligent," he said, noting that they "just ship data back raw from wherever they are." To avoid compromising the science mission, it would be better to gather data locally and have a smart, on-board system evaluate it and send a news flash to mission command on Earth "only when something proves interesting."

Pathfinding Users as Partners

On another topic: "We intend to stay focused on our customers -- the aeronautics science and engineering community, including NASA, industry, and academia -- and seek to fulfill their needs," Edwards said. "However, the present relationship has become lopsided: users get free time on a great system and they walk away happy. We want a much more two-way exchange in research collaboration."

Edwards envisions reducing the number of users (currently 1,500) and increasing the size of allocations to projects that have been carefully screened as pathfinding partners for ACNS. After the new high-speed processor system (HSP4) is available in 1998, it "should deliver processing power significantly greater than vonneumann," he added.

"We'll always push the envelope of capability and when we succeed, we'll be the premier supercomputing center for a few brief moments, until that resource gets moved to production mode," Edwards said. As the performance of the system improves, the user base will be expanded. "When the platform becomes 'production-ready,' ownership will probably be transferred (along with its customer base) to CoSMO, to be operated as a production resource."

Investing 'On the Margin'

On the state of federal budgets and the competitive pressures to become a bigger facility, Edwards observed: "We can't afford to be the biggest supercomputer center in the country. We cannot outspend other players like the Department of Energy or Department of Defense. We have to invest wisely, so the strategy is to invest 'on the margin' -- that is, do what everyone else is *not* doing. We'll try to maintain a performance niche for NASA's flight missions. That may mean developing a system with vonneumann's capability that fits into a cassette tape recorder."

"We envision creating a supercomputing power grid, or metacenter," he continued. "This is needed because demand for peak supercomputing performance usually is highly variable, with heavy use during the design phase tapering off quite a bit by the production phase. Buying the in-house capacity to satisfy peak requirements is a waste of resources. However, buying a system that supplies only average load performance leaves you with no peak capability and is equally ineffective."

"While industry users have these variable requirements, you get a nice blend that averages out when you lay them together," Edwards said. "With a supercomputer infrastructure that's available to everyone, including government, you'll have a much more efficient system overall."

He explained that the idea of a power grid is to be able to submit a job to a supercomputing service provider without worrying about details, such as what computer will run the job, where the job will be run, or how the results will be sent. "All of that will be transparent in the process."

Edwards sees this supercomputing power grid as completely heterogeneous and geographically distributed. Admittedly, "getting to that level of capability is a huge research project," he said. "The idea of transparent supercomputing capacity should drive a lot of the work we're doing here." (*See related information in* "Changes Made to CRAY J90 System.")

ACNS will focus on the research needed to make the power grid a reality: What are the system functionalities required? Can we build those or get them built? The organization is well-positioned to impact this technology area because of the NAS track record in establishing standards (such as UNIX), integrating systems, and in developing software tools (such as the Portable Batch System) that are emerging standards. "These are the kinds of things we'll do again on a bigger scale to realize this power grid," he said.

NAS Now Part of NASA Information Technology Program

The NAS Program, formerly a "focused" program within NASA's Office of Aeronautics, was merged October 1 into the Information Technology (IT) Program. IT is a 'base research and technology program" with research projects involving the efficiency and affordability of airplanes -- as well as the safety and capacity of the airspace -- at Ames, Lewis, and Dryden Flight Research Centers.

The officially named Advanced Computing, Networks, and Storage (ACNS) Program continues to be associated with the NAS Systems Division (headed by Marisa Chancellor), which serves as a matrix organization for three NASA aeronautics programs. Each program has separate funding, management, and mission goals.

Aside from ACNS, which will continue work on cutting-edge integrated systems capabilities, the division includes part of the Computational Aerosciences (CAS) portion of the High Performance Computing and Communications (HPCC) Program, which focuses on networks and multidisciplinary supercomputing. ACNS and HPCC will collaborate closely in areas where their objectives overlap.

One example of collaboration between ACNS and HPCC/CAS is expected to be the FY98 purchase and configuration of the latest generation supercomputing testbed, known as HSP4. Another example of collaboration involves combining NASA's wide area networks under one ATM system, with HPCC and ACNS each receiving some of the bandwidth, according to Tom Edwards, program manager for Information Technology (IT).

Production computing is being turned over in stages to the newly formed Consolidated Supercomputing Management Office (CoSMO) until the organization becomes fully operational in about 18 months. In addition, ACNS and HPCC will provide development efforts for CoSMO so that its computing resources continue to deliver computer cycles at the most cost-effective price available.

to the article

Changes Made To CRAY J90 System

by Elisabeth Wechsler

In order to meet computing requirements for the Aeronautics Consolidated Supercomputing Facility (ACSF), and in keeping with NASA-wide consolidation efforts, the NAS Facility's CRAY J90 computer cluster is being upgraded to increase memory, CPU, and disk capacity. The resulting design will feature a master-slave configuration: one J90 will act as primary host, with home file systems for users, the master Portable Batch System (PBS) scheduler, and all interactive work.

The other three J90s from the original "newton" system will be dedicated to batch processing only. Using the NAS-developed PBS, a set of queues will be created that allow users to submit jobs and have the scheduler assign them to the appropriate system, based on requested resources and system utilization.

"We'll also be taking advantage of features of the NAS distributed accounting package to enable a single account to be charged -- whether jobs are run on newton or eagle (the ACSF's CRAY C90)," said Bob Ciotti, newton project lead.

Workaround Required

The J90s, delivered in February 1996 and originally planned for use as a parallel cluster testbed (reported in NAS News, March-April '96), are being reconfigured to run as a production system. In September '96 planned upgrades were cancelled due to delayed availability dates for critical hardware and software.

Without those upgrades, clustering the systems became impractical, explained Bob Ciotti, newton project lead. "There was risk associated with the project to cluster these systems; however, we knew that falling back to a traditional production environment was a good risk mitigator," he said.

Benefits in Usability, Cost Savings

"The benefit to users will be a system that's as easy to use as eagle is today," noted Alan Powers, high speed processor group lead. In addition, the new configuration's focus will be on production computing, rather than clustering to solve much larger problems.

"Incorporating this system into the ACSF will roughly double the facility's available processing capacity at substantially less cost," Ciotti said. "It also adds six times (or 1.5 gigawords) the amount of available memory -- typically the most valuable resource in running large-scale scientific applications."

Production Metacenter

The change, to be implemented this spring, represents the next step toward the NAS Systems Division's goal of building an integrated computing environment, sometimes described as a power grid or metacenter (an idea discussed in the recent interview with Tom Edwards, program manager for Information Technology). "We're taking software technology developed on the testbed systems and putting it into general production. This is the first time we've operated the metacenter in a true production environment," he continued.

Development of this integrated computing environment will require several pieces of software working together to provide a single user interface into a collection of production resources. "We want to set up an environment that no matter what location, type of hardware, or time of day, users can submit jobs from their desktop workstation to run on any number of systems," Ciotti said.

Extended User Base

Once the system passes staff testing, a significant percentage of the J90 allocation will be turned over to the ACSF and managed as a production resource. Some 300 researchers, including all current ACSF account holders, are expected to use the system. They'll be asked to give input to NAS developers about adequacy of turnaround time, queue limits, and viability of applications, Powers said. "We want feedback on an ongoing basis to ensure that we're meeting users' requirements."

Researchers with projects that have been approved for access to the newton cluster will have access to the new configuration through the FY97 operational period, ending September 30. Both ACSF and NAS Facility users will share access through this period, said Chuck Niggley, scientific consulting group lead. However, the configuration will no longer be operated as a cluster, internode communication of parallel codes will not be supported, and interactive parallel jobs will no longer be permitted, he added.

New Software Tool Developed For Interactive Analysis of Wind Tunnel Experiment Data

by Samuel P. Uselton

A new software tool that displays data collected during wind tunnel tests and supports interactive analysis of this data has been developed by the NAS Systems Division's data analysis group. The tool, called exVis, is the first step toward developing a multi-source visualization environment. The ability for researchers and engineers to display and interact with data from a variety sources within a single visualization environment will increase the understanding gained from the data while reducing the time required for this process.

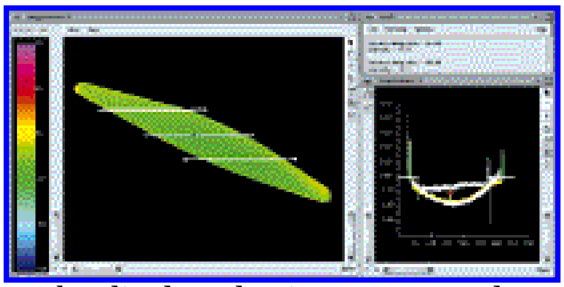

The initial prototype was developed for data collected using a pressure sensitive paint (PSP) system. The exVis software package reads a Flexible Image Transport System (FITS) format file produced by the PSP data acquisition system, which contains the coefficient of pressure (Cp) values, and then displays this data as an image in an Imageviewer" window (at left in figure). The image can be displayed using a gray scale or a pseudocolor mapping of data to screen intensity.

Data is Selected, Graphed For Analysis

Users can create tools that allow the selection of a "slice" of data from the image along a line segment. This slice is then graphed in a separate window, called "GraphViewer" (at right in above figure) that shows how the Cp varies along the length of the slice. Several slices can be selected from the same image and graphed on the same axes, or different slices (or sets of slices) can be graphed in separate windows.

Multiple ImageViewers can each display an image, with selected slices graphed together or separately. Both the Imageviewer and GraphViewer can be queried interactively to show the data value at the cursor; the GraphViewer also displays the corresponding location in the ImageViewer. Additional features include:

- size and position manipulations for both image and graph viewers
- the ability to delete a previously selected slice or graph
- the ability to print a graph produced by exVis

- saving and restoring the state of a user's session
- an online help feature

Cooperation Between Groups 'Crucial'

The exVis project began as an exploration of possibilities for collaboration between the NASA Ames data analysis group. Although exVis was developed specifically with PSP data acquisition in mind, it is sufficiently robust to work with other scalar data that is provided in FITS format, and could be easily adapted for data provided in other two-dimensional array formats.

The exVis software was integrated with other software developed for DARWIN for use in a wind tunnel test held last September. Members of the data analysis group primarily responsible for this work include Glenn Deardorff, Yinsyi Hung, Leslie Keely, and Arsi Vaziri. The cooperation of James Bell, who developed the PSP system, and other DARWIN team members including David Korsmeyer, John Schreiner, and Joan Walton, was crucial to developing a useful system.

Extensions to exVis in the Works

The exVis package is currently being extended to operate on several other types of data acquired during wind tunnel experiments and to incorporate more visualization and interaction techniques. Additional datatypes will include Doppler Global Velocimetry data and data acquired via a phased acoustic array. New interactive analysis functions include color map editing and the ability to compare data by displaying several images side by side or in an animated sequence.

This new version will be integrated with software to be used in the Aeronautics Design and Test Environment (AD/TE), an Ames-wide demonstration project in integrated system design.

A prototype second version, designated exVis 1.5, includes some of these extensions and is planned for limited release this month. The complete multi-source version, designated exVis 2.0, is planned for release sometime over the summer. A companion product, Visual Integration of Simulated and Observed Results (VISOR), is being designed to extend the methods developed in exVis to display data in a three-dimensional context and to allow interaction with scenes containing both experimental and computational data. VISOR will also be used in the AD/TE project, with a planned completion date of September.

A brief description of the initial release of exVis contains links to a sample image, a sample version of the program that can be downloaded for evaluation, and a form for ordering the software.

Samuel P. Uselton, exVis project team leader, has been developing visualization technology at the NAS Facility since 1989. He has worked in computer graphics and scientific visualization for 21 years. In addition to the aerospace data used at NAS, he has supported visualization of medical data, and oil exploration and production data.

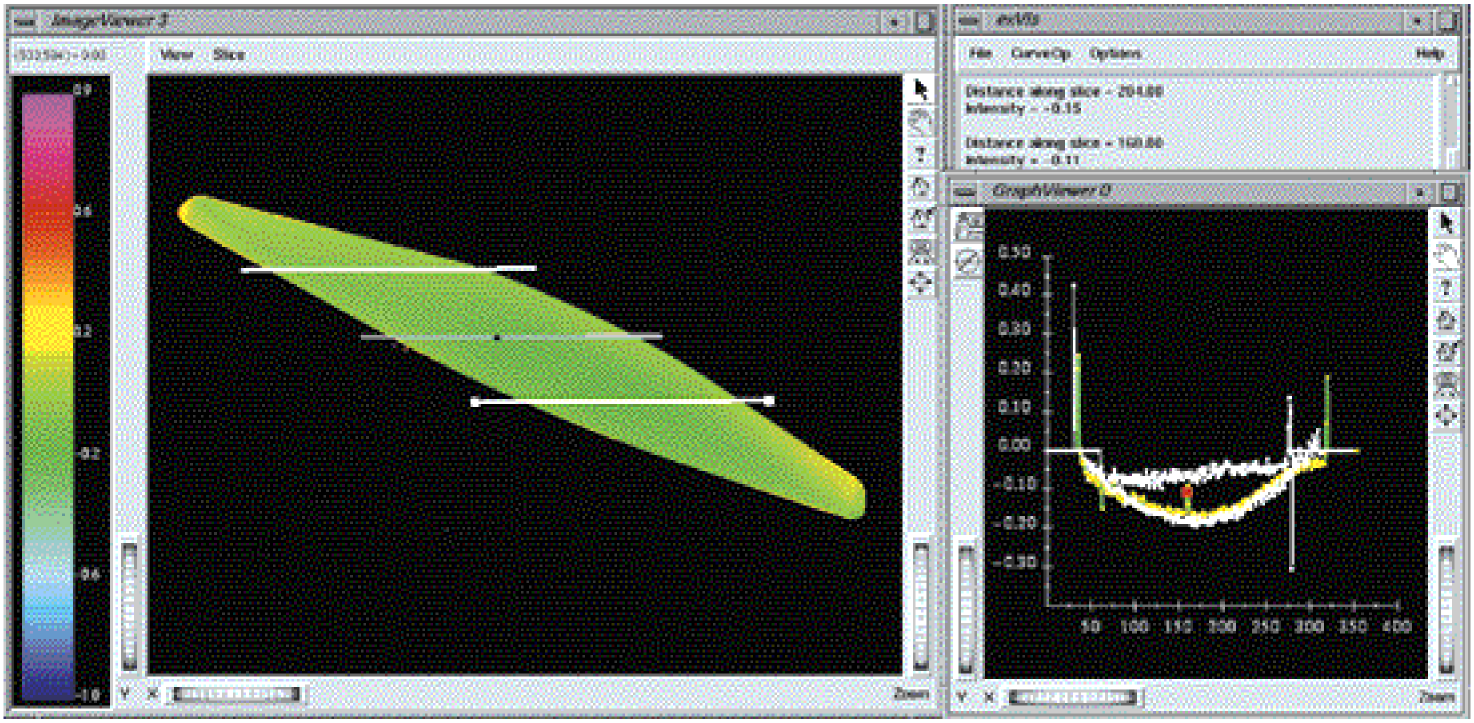

The exVis ImageViewer window (left) shows pressure-sensitive-paint data for the OAW-3 oblique all-wing model, acquired in the 9 x 7-foot supersonic wind tunnel at NASA Ames Research Center. The white line segments are slice tools that have been interactively placed on the image. The GraphViewer (right) shows x-y plots of the coefficient of pressure versus length along the tool for each selection. The highlighted peak (the red mark on the graph) was selected in the GraphViewer, causing the corresponding Imageviewer location to appear as the black spot along the gray tool. The third window (top right) is the simple user interface used to access datafiles and create ImageViewers and GraphViewers.

OAW-3 image courtesy of Robert Kennelly (NASA Ames), principle investigator for the Oblique All-Wing Project, a cooperative effort with Boeing Co. and McDonnell Douglas Corp.

to the article

High-speed Processor Techniques:

New Highbatch Queue on the CRAY C90s

by Terry Nelson

A new batch queue, called "highbatch," has been created on the NAS and ACSF high-speed processors, vonneumann and eagle (both CRAY C90 systems). This was done in response to several users who had commented that while the memory limit in the debug queue was fine, it would be helpful if the time limit were a little longer. Specifically, more time was needed to generate grids prior to job runs where five minutes was not sufficient time.

It was decided that, rather than increasing the limits on the debug queue -- which would affect the queue balance and system throughput -- it would be preferable to create a separate queue expressly for this purpose.

Queue Limits and Usage

The limits for the newly created highbatch queue are as follows: 300 megawords (MW) of memory on vonneumann and 64 MW on eagle; and an 1800-second time limit -- much less than the batch queue -- but six times more than the debug queue's time limit.

Users request the highbatch queue in a fashion similar to that used for jobs being placed in the `mt' queue. The script should contain a line such as:

```
#PBS -q high
```

along with the *-l mem* request and the *-l cput* request, not to exceed the above queue limits.

Because jobs in the highbatch queue are given priority over comparable jobs in the batch queue, they have an advantage in faster turnaround, so some control was deemed advisable. Therefore, users are allowed to run up to 1800 seconds' worth of jobs in this queue, twice a week. The 1800-second allotment is reduced by the amount used by jobs run in this queue, and is reset twice a week, at 12:05 a.m. each Monday and Thursday. Users who exceed their allotments for jobs to the highbatch queue will receive an email message stating, "Insufficient high priority allocation."

Utility Tracks Time Used

A utility called *hsee* tells users how much highbatch time they have used in the current half-week. For more information on this utility, see the man page for *hsee*. If you have any questions on using the highbatch queue, send email to **tnelson@nas.nasa.gov**.

Filesystem Reconfiguration Makes Sense for ACSF Users

by Jill Dunbar

In a move to give users of the Aeronautics Consolidated Supercomputing Facility (ACSF) more flexibility in managing their home filesystems, the NAS Facility's high speed processor (HSP) group is reconfiguring the current filesystem. The change, a process which is expected to continue through June, will also make the system easier to use, noted Daniel DePauk, HSP systems analyst.

DePauk explained that users are sometimes confused with the current two-filesystem setup on the ACSF CRAY C90 (eagle), consisting of home directories for storing small files and source code, and a directory in the Central Storage Facility (CSF) for housing large datasets for the long term.

The new filesystem structure will mirror that already in place on vonneumann, the NAS Facility's CRAY C90: a "superhome" hybrid filesystem that allows small files -- as well as large files needed in the short term -- to remain accessible, or "online." Files headed for long-term storage will be stored in the NAStore mass storage system.

The decision to convert to a superhome system ties in with upgrades to disk storage hardware, which should take place within about six months, DePauk said.

Users Need To Know:

He emphasized that users need to be aware of three items:

- All home filesystems have already been configured as the superhome filesystem.
- Users are encouraged -- and have been since last fall, DePauk noted -- to move their current working files that now reside on the CSF, to their home directories on eagle.
- Rarely accessed files should be copied from the CSF to NAStore for long-term storage; then, these files and unneeded files should be removed from eagle.

"It's preferable for users to move their own files from CSF as soon as possible," DePauk said. Otherwise, the HSP group will move these files to users' home directories on scott (one of the NAStore machines) in a directory named eagle_csf. "After that, the files will no longer be accessible by users on eagle."

Alan Powers, HSP group lead, estimated that about 1.5 million files will be moved. Of these, about 90 percent -- all files of less than 1 megabyte (MB) -- will always stay online. He noted that that amount represents less than 5 percent of the data stored.

HSP group members will notify users by email or phone one or two days in advance of file moves. Those who can't remember their scott passwords can retrieve them by contacting ACSF/NAS User Services: (415) 604-4444 or (800) 331-8737, email **nashelp@nas.nasa.gov**.

Definitions, Limitations

Elaborating on the use of superhome and NAStore, DePauk stated the following guidelines:

- "Small" files are defined as those under 1 megabyte (MB). As file space is needed, those larger than 1 MB may be migrated from disk to tape by DMF, Cray's Data Migration Facility. The presence of DMF guarantees that any file under the 1-MB limit (such as source code) will always be online.

- "Short-term" storage (for larger files used by current projects) is loosely defined as a maximum of one month.

- Unused files that need to be retained for long periods should be moved to NAStore and deleted from eagle.

In addition, each user receives a generous disk quota of 10 thousand files and 1 gigabyte of file space.

Batch Jobs Get 'Special' Treatment

For batch jobs, users are cautioned to not use their home directories as the current working directory because larger files have poor I/O performance in home directories, DePauk explained. Users will get better job turnaround if they use the following alternate procedure:

1. Copy the file to $BIGDIR (as the current working directory.
2. Change directories to $BIGDIR.
3. Execute the program.
4. After the jobs complete, copy needed files back to the home directory.

See information on the progress of the ACSF filesystem migration. or send email to hsp-snp@nas.nasa.gov.

'Commodity' Cluster Prototype Ready for Kickoff

The NAS Facility is planning a project to develop scalable systems software capable of supporting a scientific workload on thousands of processors -- a key step toward petaflops computing. "Petaflops performance will require thousands -- to hundreds of thousands -- of processors," said project lead Bill Nitzberg. "Scaling systems software to this unprecedented level requires a prototype system with commensurate scale in both number of processors and workload."

"Developing truly scalable systems software requires not only a highly parallel system, but also a real workload, which is critical to the success of any systems software development project," he continued.

The goal of the "Whitney" project consists of collecting, porting, and developing the systems software and related technology necessary to run the cluster of systems as a production testbed. The plan is to build a prototype system of 20-30 nodes, each with 1-2 processors, and begin developing components such as compilers, message-passing libraries, resource management facilities, and system administration tools).

Emphasis on Scientific Workloads

By leveraging commodity computing technology -- acknowledged to be cheap, fast, interchangeable, and readily available -- the NAS Facility will build a cluster of 200-500 commodity processors designed to run scientific workloads by FY98.

The idea of commodity PCs started to interest scientists last year when performance for 1 node achieved 25 megaflops on applications of interest. By comparison, 1 node on the IBM SP2 achieves performance of 65 megaflops. "Only in the last six months have inexpensive commodity processors become powerful enough to warrant their use in a development platform," Nitzberg said. "This means we can achieve 30 percent of the SP2's per-node performance at significantly lower cost per node," Nitzberg said.

'Non-trivial' Processing Power

The Whitney project is broken into roughly three phases: evaluating components, developing systems software, and running a production testbed. The first phase has largely been completed by groups within the NAS Systems Division and at other research organizations.

Similar projects in this area of research -- such as Beowulf at Goddard Space Flight Center, Network Of Workstations at UC Berkeley, and Real World Computing Project in Japan -- have demonstrated that a "pile of PCs" can provide nontrivial processing power at very low cost. "We believe it's now possible to support the NAS testbed workload for about $1 million," Nitzberg said.

"Processing power is only one component of a viable production capability," he noted. This project (named after cotton gin inventor Eli Whitney, who conceptualized systems with interchangeable parts) "extends previous work by attacking the larger problem of running real scientific workloads on this type of platform." And without the unpredictable -- and sometimes hostile environment of a real workload, software cannot be adequately tested, he said.

Unique Focus

Nitzberg emphasized that the NAS Facility's approach will be unique in its focus: "We aren't just reinventing the wheel. We want to take the concept further than other research centers who have demonstrated that linking lots of PCs together is feasible," he said, and "attack the larger problem or running real scientific workloads on this type of platform."

"We're doing all the system architecture ourselves and we're responsible for the system's overall balance (including making cost-benefit tradeoffs)," he said. One of the project team's first decisions will be to choose between two vastly different network technologies: 100-bit/second Ethernet -- much cheaper -- or Myrinet, fast proprietary networking hardware for clusters.

"Since the same engineering solutions may not work for the larger-scale system, we may ultimately have to trash the 20-30 unit prototype in order to design the optimum architecture for the production cluster," Nitzberg said, emphasizing that this learning curve is unavoidable.

'Very Aggressive' Goals

Toward the end of FY98, the project's goal is to deliver as much as "10 gigaflops of real, sustained application performance on this testbed," he continued, adding that this goal is "very aggressive." The NAS Facility's CRAY C90 currently delivers only 4-5 gigaflops of sustained performance on the production workload.

The final phase, running a system like Whitney as a real production platform, "hasn't yet been attempted by any other group," he said. "Since vendors are typically 6-18 months late in delivering parallel processing systems, we'd have significant performance improvement a year earlier (than vendors) if we can complete this project."

For more information on the Whitney project, send email to **nitzberg@nas.nasa.gov**.

Lewis Aeroacoustic Study Reveals Jet Noise Patterns

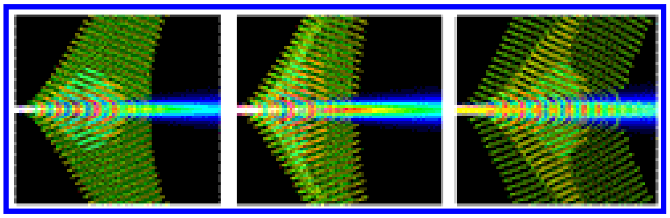

The figure at left shows the mean velocities of three jets with a snapshot of the instability pressure contours superimposed to show their noise radiation patterns. This aeroacoustic study, submitted as a NAS Technical Summary for the 1995-96 extended operational period, found that the Inverted Velocity Profile jet (right) was 6 decibels (dB) louder than the single round jet (left), while the Normal Velocity Profile jet (center) was 2 dB quieter. Calculations were performed on the NAS Facility's CRAY C90 using 5.5 megawords of memory and 4 hours of CPU time for each case.

'Catching' Large-scale Instabilities

In a supersonic jet, there are two mechanisms to generate noise: small-scale turbulence in the exhaust (which produce very small eddies) and large-scale instabilities (which make the entire jet wiggle around -- much like rising smoke that initially rises smoothly, then curls around and dissipates). "In our work, we can catch the large-scale instabilities but not the small-scale ones," said Ray Hixon, principal investigator for the project at Lewis Research Center.

Large-scale instabilities are thought to be the dominant noise producers, he explained. Instabilities for a given jet can be calculated by solving the unsteady linearized Euler equations using a high-accuracy numerical algorithm, such as the new MacCormack-type solver developed for this work. The results will aid in the design of environmentally friendly and efficient jet engines.

Inner vs. Outer Jet

By following the exhaust path of the jet in reverse towards the engine, you can observe that the inner jet is a circle and the outer jet forms a larger circle around the inner jet. The Normal Velocity Profile coannular jet has a fast inner jet and a slow outer jet; the Inverted Velocity Profile coannular jet has a fast outer jet and a slow inner jet; and the reference single round jet has an inner jet only. In this comparison, the jets have the same thrust (or force that the jet exerts) and the same amount of exhaust emitted (or, mass flow rate).

The linearized Euler equations solved in this research assume that the sound waves and other disturbances are small compared to the pressures, velocities, and so forth, created by the jet exhaust, "but the disturbances are too small to affect either the mean flow quantities or the other disturbances," Hixon said. The mean flow of the jet exhaust is calculated beforehand using an off-the-shelf numerical code, and its input to this code (density, velocities, and pressure).

3-D Extension Planned

Future research will extend the linearized Euler code to 3-D generalized coordinates and will test the code with elliptical and rectangular jet geometries. In addition, the existing code will be used for more parametric studies, Hixon said.

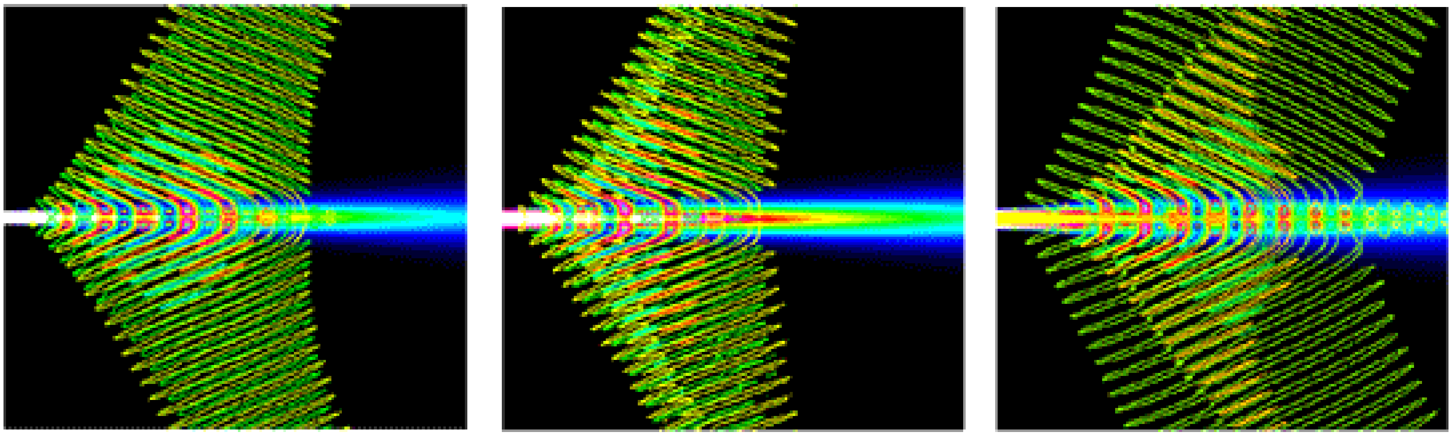

UIG Meeting Brings Users Up To Date

The annual meeting of the NAS Systems Division User Interface Group, held on January 10, showcased accomplishments for the past year at the NAS Facility and NASA Ames Research Center, and presented a preview of changes to come. Some 70 attendees, including key representatives from industry, academia, Department of Defense, and NASA centers attended the all-day event.

The overall response from participants was positive and indicated that "NAS is moving in the right directions," said Lisa Reid, customer communications group lead and UIG meeting coordinator. "The NAS UIG meeting is one way to show our ongoing commitment to maintain effective two-way communication with the user community and to be responsive to their needs," she added.

In opening remarks, Ames Center Director Henry McDonald explained changes implemented by the Office of Aeronautics regarding project approval and administration in the larger context of federal budget cutbacks. He also updated the audience on the progress of the (Ames) Center of Excellence for Information Technologies, which will result in restructuring the former NAS Program to be part of a "base research and technology" program.

Marisa Chancellor, Acting Chief of the NAS Systems Division, reviewed the Facility's accomplishments in the last year, including changes in the allocation process, staffing levels, and a long list of technical achievements, including the CRAY J90 cluster installation, new source code versions for five NAS Parallel Benchmarks, and the IBM SP2 parallel metacenter established with Langley Research Center.

NASA Information Technology (IT) Program Manager Tom Edwards spoke about the new directions and emphasis for the former NAS Program, now a component of IT, and requested audience feedback.

Several online demonstrations of visualization technology that were shown at Supercomputing '96, as well as a tour of the NAS Facility, were offered during the day. The demonstrations included the Stereo Visualization Theater, Virtual Wind Tunnel/Virtual Workbench, NAS Kiosk (a Web-based "tour" of NAS Systems Division work), and NAS Technical Summaries for FY96. In addition, a session reviewing the FY97 operational year allocation process for the NAS Faciity's CRAY C90 was led by Bryan Juliano, local area networks group lead.

Four afternoon sessions were held by NAS technical experts on the topics of testbeds, future computing architectures, integrated design systems, and state-of-the-art visualization techniques. A wrap-up session led by Edwards included summaries of these sessions.

A full report of the UIG meeting is available.

March - April 1997
Vol. 2, No. 23

Executive Editor: Marisa Chancellor

Editor: Jill Dunbar

Senior Writer: Elisabeth Wechsler

Contributing Writers: David H. Bailey, David Kenwright, Terry Nelson, Samuel P. Uselton

Image Coordinator: Chris Gong

Other Contributors: Richard Anderson, Jeff Becker, Chris Buchanan, Nicholas Cardo, Bob Ciotti, Daniel DePauk, James Donald, Tom Edwards, Pat Elson, Michael Gerald-Yamasaki, Ray Hixon, Leslie Keely, Tim Kirby, Chris Kleiber, Chuck Niggley, Bill Nitzberg, John Parks, Alan Powers, Marcia Redmond, Lisa Reid, Ken Stevens, Leigh Ann Tanner, Dani Thompson

Editorial Board: Nicholas Cardo, Marisa Chancellor, Jill Dunbar, Chris Gong, Mary Hultquist, David Kao, Chuck Niggley, Elisabeth Wechsler

NAS NEWS

Volume 2, Number 23

March – April 1997

Onward to Petaflops: Supercomputing for the 21st Century

by David H. Bailey

Continued on page 6

THIS ISSUE

New "NAS" Program Manager Speaks Out
page 3

exVis Tool for Interactive Analysis
page 5

Changes to CRAY J90 Cluster
page 4

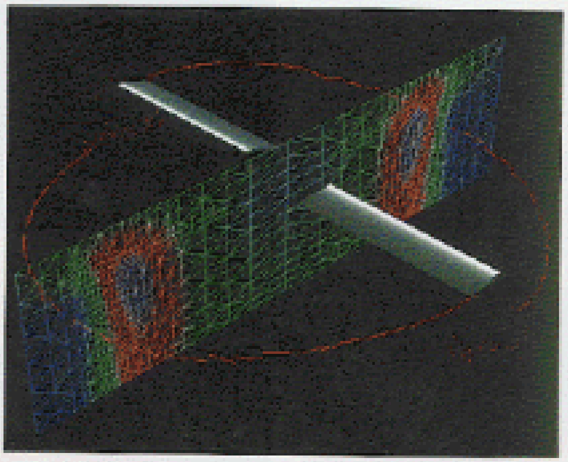

Vortex Core Detection Technique Helps Rotorcraft Research

by David Kenwright

Problem With Diffusion

Vortex Cores Steer Grid Generation

Continued on page 2